

Towards Unstructured Mesh Generation Using the Inverse Poisson Problem

Guy Bunin

February 17, 2008

Abstract

A novel approach to unstructured quadrilateral mesh generation for planar domains is presented. Away from irregular vertices, the resulting meshes have the properties of nearly conformal grids. The technique is based on a theoretical relation between the present problem, and the inverse Poisson (IP) problem with point sources. An IP algorithm is described, which constructs a point-source distribution, whose sources correspond to the irregular vertices of the mesh. Both the background theory and the IP algorithm address the global nature of the mesh generation problem. The IP algorithm is incorporated in a complete mesh generation scheme, which also includes an algorithm for creating the final mesh. Example results are presented and discussed.

1 Introduction

Boundary alignment is a critical feature of meshes in many applications. In a boundary aligned mesh the boundary, or some other line, is traced by the sides of high-quality cells, see fig. 1. The definition of a “well-shaped” cell may be application dependent, but in many cases, cells similar in shape to squares (for quadrilateral cells) or equilateral triangles (for triangles) are preferred. Characteristics of the entire mesh are also important, such as smooth cells-size and cell-shape transitions.

The problem of producing boundary aligned meshes with well-shaped cells has been the subject of extensive research [1]. Still, many popular algorithms are heuristic in nature, and a more general understanding of the subject is called for, especially when quadrilateral meshes are considered. A key difficulty is the problem’s *global* character: the shape and position of every cell in the mesh is, at least in principle, related to that of any other cell.

In a previous work [2], we described a relation between the problem of two-dimensional unstructured mesh generation, on both planar and curved surfaces, and another well-known problem, namely the *Inverse Poisson* (IP) problem. The IP problem is concerned with reconstructing a *source distribution* ρ of the Poisson equation $\nabla^2 \phi = \rho$, from information on the potential ϕ at the boundaries. In that work, the mesh was assumed to be *conformal* away from the irregular vertices (vertices whose degree is different than four), like a grid mapped by a conformal mapping. Such grids have the property of having square cells in the limit of an increasingly finer grid. Under this assumption, the problem of mesh




Figure 1: Part of a boundary aligned mesh (black lines). The gray line represents the domain’s boundary.

generation was then shown to reduce to an IP problem. The irregular vertices of the mesh correspond to point sources (delta functions) of ρ , and ϕ is interpreted as the logarithm of the local resolution. This theoretical framework turns the focus to the irregular vertices of the mesh: once their distribution is fixed, the continuum properties of the mesh - local resolution and directionality - are known. Note that it is also an explicitly global formulation, since the resolution at any given point is affected, via the function ϕ , by the locations of all irregular vertices in the mesh.

In this paper the generation of planar quadrilateral meshes is discussed. Resting on the results of [2], a new IP algorithm is presented, designed to construct source distributions of the appropriate type, which approximate the resolution and mesh directionality inputs at user-specified points, such as at the boundaries. The IP algorithm is then incorporated into a complete mesh generation scheme, which also includes a technique for generating the final mesh. An implementation is described, and shown in example cases to generate boundary aligned meshes, where well-placed point sources create smooth cell transitions and high quality cells. A similar procedure is probably applicable to triangular meshes, but is not discussed in the present work.

Remeshing of curved surfaces has recently attracted considerable attention; for a review see [3]. Many of the algorithms receive an input mesh directionality throughout the surface, usually the principal curvature directions of the surface. This setting presents different challenges than those addressed here, since the mesh structure is determined, to a large extent, when mesh directionality is given everywhere in the domain. For example, the locations of critical irregular vertices are dictated by the mesh directionality in the *vicinity* of these points¹. Another related subject is surface parameterization, concerned with creating mapping of surfaces to the plane. Conformal surface parameterizations are created in [8], but boundary alignment is not addressed.

The paper is organized as follows: Section 2 shortly reviews the relevant theoretical background, with emphasis on the relation between the IP problem and unstructured mesh generation. Sections 3 and 4 describe the proposed IP algorithm, and the mesh generation scheme. Section 5 describes an implementation of the algorithm, and section 6 gives examples of meshes generated. Conclusions and possible directions for future research are discussed in section 7.

2 Background theory, relation to the IP problem

In this section an overview of the background theory will be given. Section 2.1 explains the rationale underlying the mathematical formulation. In section 2.2, some key conformal geometry concepts are discussed, together with their relevance to the present problem. Section 2.3 summarizes the results developed in [2], relating mesh generation with the IP problem. The exposition is limited to planar two-dimensional mesh generation; A detailed account, in the more general setting of curved surfaces, can be found in [2].

2.1 Motivation

We start by considering mesh generation using conformal mappings, which can be viewed as a special case of the theory to be described. Mapping techniques, in general, construct a function from one domain, for which a mesh already exists, to a second domain, which is being meshed. The mapping function is then used to map the mesh into the second domain. A key idea is that continuum properties of the mapping function control the shapes of the cells of the new mesh, at least for small enough cells. For example, if the mapping is conformal, i.e. angle preserving, a cell with right inner angles (rectangle) will be mapped to a target cell with approximately right inner angles.

In unstructured mesh generation the connectivity of the mesh is not known in advance, and a more general framework is called for. In what follows, the interplay between the two domains which serves as a paradigm of mapping techniques, is replaced with an interplay between two definitions of distances on the input domain. Instead of imagining the mesh as being the image of a mesh on a different domain,

¹In many works, *both* mesh directionality and local resolution are specified on the entire surface, and since conformality requires a specific relation between the two properties, the resulting meshes are not conformal (nor are they claimed to be). In one exception [4], a preprocessing step attempts to adjust the local resolution, in order to create a more closely conformal mesh.




Figure 2: **a.** A conformal grid. Solid lines are geodesics. The distance between two parallel lines is \tilde{h}_0 , measured with the new distance definition. Geodesics bend towards smaller cells. Crosses represent the cross-field at selected points. **b.** The curvature of a geodesic is related to the gradient of ϕ .

we *redefine the distances* on the domain to be meshed, and *fix the cell edge length*. Thus, the local mesh resolution will be proportional to the new local distance definition: a large new distance between two given points will mean more cells will be placed in between, hence a higher local resolution. Distances are redefined using the concept of a *metric*, known from Riemannian geometry.

Since we focus on cells which are squares in the limit of an increasingly finer mesh, we need to define only two local properties: the resolution (inverse of cell-size), and the direction of the cell-edges. The new resolution, as noted before, is controlled by a new distance definition. We use a new distance which is locally a scaling of the old distances: the new distance between a point and other nearby points is equal to the old distance, multiplied by a scalar factor which is independent of the direction. Thus, a small square measured in one distance definition is also (approximately) a square according to the new distance definition. In Riemannian geometry terminology, the old and new metrics are said to be *conformally related*. The mesh directionality is related to the resolution, as is described in the following section.

2.2 Definitions, conformal geometry relations

It is convenient to work with a function $\phi(\vec{r})$, defined as the *logarithm* of the local scaling factor. That is, a small square of side length h measured with the original, Cartesian distance definition, is a square with side length $\tilde{h}(\vec{r}) = h(\vec{r}) e^{\phi(\vec{r})}$ as measured with the new distance definition, see fig. 2,a.. If we imagine that the domain is covered with many small squares, all with the same side length \tilde{h}_0 in the new distance definition, the size of these cells in the original distance definition will be $h(\vec{r}) = \tilde{h}_0 e^{-\phi(\vec{r})}$. The (local) resolution is the inverse of the local-cell-size, hence

$$resolution \propto e^{\phi(\vec{r})}, \quad (1)$$

where the proportionality constant is a single number for the whole mesh.

The second continuum property is the local directionality of the cell edges. The edges should ideally meet at right angles everywhere, except at *irregular vertices*, where the number of edges incident on the vertex is different than four. It is therefore natural to assign to every point a set of four directions, mutually parallel or perpendicular. This concept was expressed by many authors, and given various names, such as *mesh directionality* [6], *4-symmetry direction field* [7], and *frame field*, although the last may refer to a structure which also holds cell-size information [5]. Graphically, this object can be represented by a cross at every point, see fig. 2,a., and here will be called a *cross-field* [2]. On the plane, the cross direction can be measured by the angle θ from the x -axis to one of the directions of the cross. This angle is fixed up to an addition of $\pi/2$ radians, i.e. θ_1, θ_2 represent the same cross iff $\theta_1 = \theta_2 + n\pi/2$, for some integer n .

The function ϕ and the cross-field are not unrelated; due to conformality, lines that trace the edge directions bend towards the side with smaller cells, see fig. 2,a.. In the continuum theory, these lines

are known as *geodesics* of the manifold. Geodesics are a generalization of the concept of straight lines to non-Euclidian geometries. In the original, Cartesian coordinates, the geodesic obeys the following differential equation:

$$\kappa = \frac{\partial \phi}{\partial n}, \quad (2)$$

where κ is the curvature of the geodesic in the Cartesian coordinate system, and $\partial \phi / \partial n$ is the derivative of the function ϕ in the direction normal to the tangent, see fig. 2,b.. Eq. (2) allows one to calculate the change in the direction of the cross-field between two points connected by a geodesic. There is also a direct way of calculating the change in the direction of a cross-field between any two points, along any curve connecting the two points, known as *parallel-transport*. Let α be some curve from point a to point b , and let θ_a, θ_b be the angles of the crosses at points a, b respectively (as noted above, the angles are defined up to an addition of a $\pi/2$ radians). Then

$$\int_{\alpha} \frac{\partial \phi}{\partial n} ds = \theta_b - \theta_a. \quad (3)$$

where the integration denotes a line integral along the curve α , according to the length parameter s on α , as measured in the original coordinate system. The differential formulation of this equation is

$$\frac{\partial \phi}{\partial n} = \frac{\partial \theta}{\partial t}, \quad (4)$$

where (\vec{t}, \vec{n}) are a pair of perpendicular vectors that form a right hand system.

Where the function ϕ is defined, that is, at any point in the domain that is not a singularity, ϕ can be shown to *harmonic*, that is to obey the Laplace equation:

$$\nabla^2 \phi \equiv \frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial y^2} = 0, \quad (5)$$

where, again, the derivatives are taken with the original coordinate system. Eq. (2,5) are well-known results of conformal geometry [9],[10]. Eq. (3) is derived in [2]

2.3 Relation to the IP problem

Equations (2,5) fully describe the relations between the cross-field and the function ϕ , at any regular point of the domain, that is, any point that is not a singularity of the function ϕ . The singularities of the function ϕ are a key ingredient of the theory, since they correspond to the irregular vertices of the mesh, and unstructured meshes are those which contain irregular vertices. A detailed analysis of the possible of singularities of the harmonic function ϕ , and their effect on the resulting mesh, was carried out in [2], and shows that the only type of singularity that corresponds to a mesh with a finite number of cells is of the type $\phi(\vec{r}) \propto \ln |\vec{r} - \vec{r}_0|$, where \vec{r}_0 is the location of the singularity. In the IP literature, such a singularity is known as a *point source*. Furthermore, the prefactor of the logarithm is directly related to the degree of the irregular vertex in the final mesh. More specifically, suppose there are n_c singularities in a domain D , at points $r_m, m = 1..n_c$, then the function ϕ can be written as

Condition 1:

$$\phi(\vec{r}) = \phi_L + \sum_{m=1}^{n_c} \frac{Q_m}{2\pi} \ln |\vec{r} - \vec{r}_m| = \phi_L + \frac{1}{4} \sum_{m=1}^{n_c} k_m \ln |\vec{r} - \vec{r}_m|, \quad (6)$$

where ϕ_L is a harmonic function, $k_m \in \mathbb{Z}$, and $k_m > -4$. The numbers Q_m are known as *charges*, so Condition 1 states that the charges are integer multiples of $\pi/2$. We will refer to the numbers k_m simply as the “ k -values” of the singularities. The degree (number of incident edges) of the irregular vertex corresponding to the source at \vec{r}_i is equal to $4 + k_m$. So, for example, irregular vertices of degrees 3 and 5 will correspond to a singularities with $k_m = -1$ and $k_m = +1$, respectively.

In a geometrical context, a logarithmic singularity of ϕ represents a *cone-point* - the tip of a cone - in the new distance definition. The charge corresponds to the angle deficit of the tip of the cone. Related




Figure 3: Conditions 3 and 4.

subjects where manifolds with cone-points are considered include the study of disclinations in elastic media [11], and surface parameterizations [12].

An unstructured mesh is required to be aligned with the boundary. Namely, that one of the cross directions be parallel to the tangent to the boundary (where a tangent exists):

Boundary Alignment definition:

$$\tan \left(\theta(p) + n \frac{\pi}{2} \right) = \frac{dy}{dx}(p), \quad p \in \Gamma, \quad (7)$$

where $(y(s), x(s))$ is a curve tracing the boundary, $p = (x, y)$ and $n \in \mathbb{Z}$. Since θ is fixed up to an addition of a multiple of $\pi/2$ radians, this definition depends only on the cross-field itself, and not on the particular choice of θ .

Requiring boundary alignment, as defined in eq. (7) can be shown to be equivalent to three more conditions on the function ϕ [2]. The following, Condition 2, applies to a point on a smooth section of the boundary, Γ . It roughly states that a smooth section of the boundary is a geodesic. This provides information on the derivative of ϕ normal to the boundary:

Condition 2:

$$\frac{\partial \phi}{\partial n}(p) = \kappa(p), \quad p \in \Gamma. \quad (8)$$

κ is the boundary curvature at p . The next condition is concerned with *junction points* of the boundary, where the tangent to the boundary is discontinuous. It assures that the cross-field will be aligned with the boundary on both sides of the junction point. Let α be a curve from point a to point b on both sides of a point c of the boundary, see fig. 3,b.. Then

Condition 3:

$$\int_{\alpha} \frac{\partial \phi}{\partial n} ds = n \frac{\pi}{2} + \theta_{in}, \quad (9)$$

where θ_{in} is the angle of the boundary, and $n \in \mathbb{Z}$. This means that for some inner angles ϕ will contain a singularity at c ; at a distance r from c , the singularity will be of type $\partial \phi / \partial r \sim 1/r$.

The final condition is concerned with the relation between two different boundary components. It is just a restatement of eq. (3) for two points on two different boundary segments.

Condition 4:

Let a, b be two points on two boundary curves Γ_a, Γ_b . Let α be a curve from a to b . Then

$$\int_{\alpha} \frac{\partial \phi}{\partial n} ds = \theta_b - \theta_a. \quad (10)$$

Conditions 1-4 form a complete set of conditions on ϕ , for a cross-field that is boundary aligned to exist. That is, conditions 1-3, and condition 4 on a number of selected curves, one from a selected point to each boundary component, are sufficient for a boundary aligned cross-field to exist, see [2].

3 Algorithm

3.1 Mesh Generation Steps

The mesh generation scheme consists of the following steps:

1. Setup and solution of the IP problem. Finds the sources' locations and charges. Described in sections 3.2-3.6.
2. Solution of the direct Poisson problem, to obtain the functions ϕ, θ throughout the domain, see section 3.7.
3. Generation of the final mesh, see section 4.

3.2 Input

Suppose that we are given an open domain D , with boundary Γ , and are given the required resolution function on Γ . Using eq. (1), this is readily translated to the required value of ϕ on the boundary:

$$\phi|_{\Gamma} = \ln(\text{resolution})|_{\Gamma}. \quad (11)$$

Condition 1 (eq. (6)) is a solution of the Poisson equation $\nabla^2 \phi = \rho$, with point sources, whose locations and charges are yet unknown. The problem is to find a distribution of sources (location and charge) adhering to the boundary alignment definition, eq. (7), or alternatively to Conditions 2-4 (eq. (8)-(10)), as well as to eq. (11). Such a problem is known as an *Inverse Poisson* (IP) problem. The IP problem may be compared with the *direct Poisson* problem, where the source distribution is given, as well as some boundary information (e.g., Dirichlet or Neumann boundary conditions), and the value of the function ϕ is to be found. In the IP problem, the source distribution is unknown, and a source distribution adhering to the known boundary information is to be found.

IP problems have important applications in various areas of science and engineering [13],[14],[15],[16],[17]. By its nature, the IP problem is ill-posed, and the solution might not be unique, and may be sensitive to small changes of the input, such as small changes in boundary conditions. In delicate problems of this type, any prior information on the source distribution may greatly affect the applicability of a specific solution procedure. In the present problem, we seek a point source distribution. A number of algorithms for solving an IP problem with point sources appear in the literature. In [18], an inverse problem where all sources have the same charge is solved. In our implementation, however, at least two different charge values must be incorporated ($Q = \pm\pi/2$). In [19],[20],[21] an inverse problem where both the locations of sources and their charges are unknown, and are reconstructed. This gives more freedom in the reconstruction than we can allow, since for the present purposes the charges must be multiples of $\pi/2$. Another important aspect of the present problem is that the domain of the IP problem may be of any shape and topology (i.e., may contain holes), whereas the above works only deal with a simply connected region (a circle, usually).

It is important to note that in the present application, in contrast to other standard applications, the input to the IP problem is not generated by some existing source distribution (perhaps with some added noise), but by the domain's shape and input resolution. The existence of a source distribution which reconstructs the input data, at least approximately, is therefore not obvious. This is an interesting and important subject, but is beyond the scope of the present work.

3.3 Complex Formulation

As is well known, the real and imaginary parts of a complex analytic function are harmonic functions [22]. This correspondence has been utilized in IP algorithms [20],[19],[21], and will be used here as well. We define the complex-valued function

$$F(z) = h(z) + \frac{1}{4} \sum_{m=1}^{n_c} k_m \ln(z - z_m), \quad (12)$$

where $z_m = x_m + iy_m$, with (x_m, y_m) the components of \vec{r}_m , that were defined in Condition 1. h is a function on D , such that

$$\text{Re}(h) = \phi_L.$$

Then, recalling that $\text{Re}(\ln(z)) = \ln|z|$, it follows that

$$\text{Re}(F(z)) = \phi(\vec{r}), \quad (13)$$

with $\phi(\vec{r})$ defined in Condition 1, and $z = x + iy$, with $\vec{r} = (x, y)$.

The functions $h(z)$ and $\ln(z - z_m)$ for some z_m are analytic in a neighborhood of any point that is not a singularity of ϕ . However, as functions over the entire domain D , they may be *multi-valued*. Multi-valued functions accept many-values at a point, depending on the path taken to that point. It is well known that the complex function $\ln(z)$ is multi-valued. Defining $\ln(z)$ as $\ln(z) \equiv \int_{\gamma} dt/t$, where γ is a path from 1 to z , the imaginary part of $\ln(z)$ is only fixed up to an addition of multiples of 2π , depending on the path taken. If the domain D is not simply-connected (i.e., contains holes), then the function $h(z)$ may also be multi-valued, since by following different paths around the holes, different values of $h(z)$ may be obtained.

The function $F(z)$, being a sum of multi-valued functions, may also be multi-valued. According to eq. (13), the real part of F is single-valued, as it is equal to ϕ at that point. We thus turn to examine the imaginary part of F . The real and imaginary parts of a complex function are related by the Cauchy-Reimann (CR) equations [22]. Recall that for a complex differentiable function $f(z) = u(z) + iv(z)$, where $u(z), v(z)$ are real-valued, the Cauchy-Reimann equations read $\partial_x u = \partial_y v$, and $\partial_y u = -\partial_x v$. This allows one to recover the imaginary part of a analytic function if its real part is given. Writing eq. (4) in the two right hand coordinate systems $(x, y), (y, -x)$, gives the relations $\partial\phi/\partial y = \partial\theta/\partial x$ and $\partial\phi/\partial x = -\partial\theta/\partial y$. These are precisely the CR equations for the complex function $\phi - i(\theta + C)$, where C is any real constant, hence

$$F(z) = \phi(z) - i \int_{\gamma} \frac{\partial\phi}{\partial n} ds = \phi(z) - i(\theta(z) + C). \quad (14)$$

The constant C is arbitrary, and will be taken to be zero. The multi-valued nature of F is explicit in the integral formulation in eq. (14). The right-hand-side (RHS) of eq. (14) shows that since the imaginary part of F is $-\theta(z)$, and θ must be fixed up to additions of $\pi/2$, then the integral

$$\int_{\gamma} \frac{\partial\phi}{\partial n} ds \quad (15)$$

must also be fixed up to additions of $\pi/2$. Indeed, this was shown to hold if Conditions 1-3 (eq. (6)-(9)) hold, see [2].

To summarize, the problem is restated as finding a multi-valued complex function F , given by eq. (12), and whose value at points on the boundary of D is:

$$F(z) = \phi(z) - i\theta(z). \quad (16)$$

In the following sections, an algorithm for solving the problem as it was here restated is described. The IP algorithm first removes the contribution of $h(z)$ from the boundary conditions, paying special attention to the junction points, and then constructs the source distribution.

3.4 Handling junction points

The function ϕ , and hence the function F , may be singular at junction points. Special procedures for addressing this behavior are taken as part of two different steps of the algorithm:

- a. Choosing the input resolution near junction points.
- b. Special treatment of junction points when h is removed from the boundary conditions.

The two subjects are discussed in the following subsections, 3.4.1 and 3.4.2.

3.4.1 Choosing the input resolution near junction points

Except in the special cases when θ_{in} is a multiple of $\pi/2$, the behavior of ϕ at a small neighborhood in D of a junction point is singular, with the singularity at the junction point, see Condition 3 (eq. (9)). If the input resolution in a small neighbourhood of the junction point does not match this singular behavior, solution of the IP problem will feature sources at any distance from the junction, no matter how small, resulting in a distribution with an infinite number of sources. In practice, this means that an IP algorithm will cluster many sources near the junction point, in a futile attempt to reconstruct the boundary conditions there. To avoid this problem, we adjust the *input* ϕ at a neighborhood of the junction point. This adjustment should vanish rapidly at a distance larger than 1-2 cell-sizes, so as to have little effect on the original required resolution.

Condition 3 states that the gradient of ϕ near a junction point must diverge as $\partial\phi/\partial n \sim 1/r$, where r is the distance from the junction point. Such a flux is formed by a logarithmic singularity at the junction point (see also [2], section 7.1). Consider a singular source term of the form $\phi(r) = \frac{Q_J}{2\pi} \ln|r|$, where r is the distance from the junction point. Then the flux through a circular arc α_r , at distance r from the junction point is (see also fig. 3,a.):

$$\int_{\alpha_r} \frac{\partial\phi}{\partial n} ds = -\frac{Q_J}{2\pi} \theta_{in},$$

where θ_{in} is the junction inner angle. Then by Condition 3, $\int_{\alpha_r} \frac{\partial\phi}{\partial n} ds = k\frac{\pi}{2} + \theta_{in}$, hence

$$Q_J = 2\pi \left(n_J \frac{\pi/2}{\theta_{in}} - 1 \right). \quad (17)$$

n_J is integer, which must be positive, otherwise a mesh with an infinite number of cells will result ([2], section 7.1). In fact, n_J has a simple interpretation: it is the number of cells incident upon the junction point in the resulting mesh². A reasonable choice for n_J is therefore

$$n_J = \text{round} \left(\frac{\theta_{in}}{\pi/2} \right), \quad (18)$$

for which the inner angles of the cells incident on the junction point are closest to $\pi/2$.

In order to restrict the effect of this correction of ϕ to a small region in D , a source term with the opposite charge $-Q_J$ can be placed *outside* D , at a distance of about 1-2 edge lengths. See for example the second and third examples in section 6.

3.4.2 Treatment of junctions when removing the harmonic part

In section 3.5 a method will be described for calculating and removing the contribution of $h(z)$ to the boundary conditions. At junction points, however, the function $h(z)$ may be singular: as D is an open set, boundary singularities are part of $h(z)$, so care should be taken when performing the calculations described in section 3.5, especially in a numerical implementation of the technique.

To avoid these problems, we subtract the junction singularities from the boundary value of $F(z)$ before proceeding with removing $h(z)$. Let z_j^J be the locations of the junction points, and Q_{J_i} the charges as given by eq. (17),(18). The complex function corresponding to $\ln|z|$ is $\ln(z)$, so the contribution of the junctions is

$$- \sum_i \frac{Q_{J_i}}{2\pi} \ln(z - z_j^J). \quad (19)$$

If a singularity is located on an inner boundary component, the logarithms must contain a branch cut somewhere in D . In order to avoid adding a cut, we add another term to every hole, equal to minus the

²During the creation of the final mesh connectivity (section 4) nearby cone-points may be joined, and cone-points may be shifted to the boundary, depending on the final cell-size. In such a case, the number of cells incident on a junction point, as well as on other boundary points, may change.




Figure 4: Removing the harmonic part using the Cauchy integral. a. The case of a simply-connected domain. b. A domain with more than one boundary component. Thick line represents the boundary, thin line the cut introduced.

total charge in each hole. Denote the inner boundaries by $\Gamma_2, \dots, \Gamma_{n_{bnd}}$, and let $z_2^{hole}, \dots, z_{n_{bnd}}^{hole}$ be arbitrary points inside the respective holes. The following terms are therefore added to $F(z)$:

$$F(z) \rightarrow F(z) - \sum_{J_i} \frac{Q_{J_i}}{2\pi} \ln(z - z_j^J) + \sum_{m=2}^{n_{bnd}} \sum_{J_i \in \Gamma_m} \frac{Q_{J_i}}{2\pi} \ln(z - z_m^{hole}) \quad (20)$$

In this way, the additional terms form an analytic function in D , with no branch cut, and are a part of $h(z)$. Since the contribution of $h(z)$ is removed from $F(z)$ as described in the next section, the exact form of this additional term, i.e. the choice of points z_m^{hole} , does not affect the results.

3.5 Removing the harmonic part

In this section a technique for calculating $h(z)$ is presented. After $h(z)$ is calculated, its contribution to $F(z)$ at the boundary can be subtracted. To simplify notation, let $F_p(z)$ denote the sum of logarithms in $F(z)$:

$$F_p(z) \equiv F(z) - h(z) = \frac{1}{4} \sum_{m=1}^{n_c} k_m \ln(z - z_m). \quad (21)$$

In the context of IP algorithms, the idea that the contribution of the harmonic part to the solution can be removed from the boundary information was suggested in [21]. There, an IP problem in the unit disk B is discussed. The value of the analytic part of $F(z)$ on the boundary of B can then be calculated by taking the Fourier transform of $F(z)$, and leaving only the positive frequency components, as can be shown by considering the Laurent series of $F(z)$ in the unit disk. By its construction, this technique applies to functions defined on the unit disk. It can be extended to other domains, if a conformal mapping of the domain to the unit disk (which exists according to the Riemann mapping theorem) is calculated. For our present purposes, since the domain D in every problem input is different, using this technique would require constructing a mapping to the unit disk for each domain being meshed separately. Furthermore, the domains in our problem may contain holes, which further complicates the matter. We therefore use an alternative approach, which is now described. It is based on the Cauchy integral theorem, and is similar to the “sum factorization” step in the Weiner-Hopf technique [23], applied to a bounded domain (Ω , defined below) in place of an infinite strip.

As in [21], we assume there is a positive distance d between the boundary of D and the source closest to it, as is the case, for example, when the number of sources is finite. Let Ω be the set of points at a distance smaller than d to the boundary, see fig. 4,a.. As will become apparent later, the value of d does not enter the calculation and is irrelevant, as long as there exists some $d > 0$ as required.

We first consider a simply-connected domain D , with boundary Γ . Define the flux through any curve α :

$$\Phi_\alpha \equiv \int_\alpha \frac{\partial \phi}{\partial n} ds. \quad (22)$$

The flux Φ_Γ through the boundary can be calculated: the value of $\partial\phi/\partial n$ at smooth boundary points is known from Condition 2 (eq. (8)), and the flux at junction points is calculated from Condition 3, as is explained in section 3.4.1 above. This flux can be shown to be a multiple of $\pi/2$.

Remark 1 *Using Conditions 2,3 and the fact that the rotation of a tangent in a simple curve is 2π , the total flux through a boundary component Γ_m can be shown to be equal to $s\frac{\pi}{2}(-4 - \sum_{J_i \in \Gamma_m} (n_{J_i} - 2))$, where J_i are the junction points, n_{J_i} for each junction is given by eq. (18), and $s = +1, -1$ for inner and outer boundary components, respectively. See also [2].*

To proceed with the decomposition calculation, we assume that $\Phi_\Gamma = 0$, as calculated from eq. (22) above. If this is not the case, following [21], we subtract a source term to the boundary conditions, centered at some point z_{in} inside the domain:

$$F \rightarrow F - \frac{k_{in}}{4} \ln(z - z_{in}) \quad (23)$$

such that Φ_Γ is zero. As will be shown below, the choice of z_{in} does not affect the results. Since Φ_Γ was equal to a multiple of $\pi/2$ before, $k_{in} \in \mathbb{Z}$. When $\Phi_\Gamma = 0$, the function $F(z)$ is single-valued and analytic in Ω , as follows from the integral representation of $F(z)$ in eq. (14), along with $\Phi_\Gamma = 0$. We now use the Cauchy's theorem, stating that for every point $a \in \Omega$:

$$F(a) = \frac{1}{2\pi i} \int_{\partial\Omega} \frac{F(z)}{z - a} dz. \quad (24)$$

The boundary $\partial\Omega$ of Ω has two boundary connectivity components: the outer boundary (which is also the boundary of D), and the inner boundary. Denote them by $\partial\Omega_o$ and $\partial\Omega_i$ respectively, see fig. 4,a.. Eq. (24) now reads

$$F(a) = \frac{1}{2\pi i} \int_{\partial\Omega_o} \frac{F(z)}{z - a} dz + \frac{1}{2\pi i} \int_{\partial\Omega_i} \frac{F(z)}{z - a} dz. \quad (25)$$

The function $\text{Re}(F) = \phi$ is harmonic on Ω . According to a decomposition theorem (see [24], Chapter 9) its decomposition into $\text{Re}(F_p)$ and $\text{Re}(h) = \phi_L$, is also unique. It follows that the decomposition of the multi-valued complex function F into F_p and h is unique, up to the arbitrary constant C in eq. (14), chosen before to be zero. The two integrals expressions on the RHS of eq. (25) correspond exactly to the two components of the (unique) decomposition of $F(z)$ in Ω as described in [24], hence

$$h(z) = \frac{1}{2\pi i} \int_\Gamma \frac{F(z)}{z - a} dz; \quad F_p(z) = \frac{1}{2\pi i} \int_{\partial\Omega_i} \frac{F(z)}{z - a} dz, \quad (26)$$

where we have used the fact that $\Gamma = \partial\Omega_o$. We would like to find F_p on Γ , but $\partial\Omega_i$ is not known, nor is $F(z)$ on $\partial\Omega_i$, so $F_p(z)$ cannot be computed directly from the second equation. The first equation can however be used, since $F(z)$ is given on Γ , so $h(z)$ can be calculated on Γ (more precisely, since h is defined in D , it is the limit of h as Γ is approached). Then $F_p(z)$ on Γ is given by $F_p(z) = F(z) - h(z)$, according to eq. (21).

Finally, we add back the source term subtracted before

$$F_p \rightarrow F_p + \frac{k_{in}}{4} \ln(z - z_{in}). \quad (27)$$

We now turn to the case when D is not simply-connected. We add source terms: one inside the domain, as discussed above (eq. (23)), and one inside every hole (i.e. *outside* D), such that the flux Φ_{Γ_m} through every boundary connectivity element is zero, see Fig. 4,b..

$$F_p \rightarrow F_p - \sum_{m=2}^{n_{bnd}} \frac{k_{hole_i}}{4} \ln(z - z_m^{hole}) \quad (28)$$

The z_m^{hole} can be the same points used in section 3.4.2, or other points. The results do not depend on the additional charges' locations, as will be explained below. We now introduce cuts so that the boundary

contains a single connectivity element (the cuts play a part in the derivation, but drop out of the final calculation), see fig. 4,b.. The cuts introduced also serve as the branch cuts of the logarithms. Denote this new boundary $\tilde{\Gamma}$. Using $\tilde{\Gamma}$, we proceed as when D is simply connected: define Ω as before, see fig. 4,b., and subtract the source term inside D , as in eq. (23). According to eq. (26):

$$h(z) = \frac{1}{2\pi i} \int_{\tilde{\Gamma}} \frac{F(z)}{z-a} dz = \frac{1}{2\pi i} \int_{\Gamma} \frac{F(z)}{z-a} dz + \frac{1}{2\pi i} \int_{cuts} \frac{F(z)}{z-a} dz. \quad (29)$$

The second integral on the RHS denotes the integral over the cuts introduced to form $\tilde{\Gamma}$. Note that each cut-path is traversed twice, back and forth. Since the flux Φ_{Γ_m} through each and every hole boundary is zero, The value of $F(z) = \phi - i \int \partial_n \phi ds$, when integrated along $\tilde{\Gamma}$, is continuous across the branch cuts, and the integrations over each cut traversed in both directions cancel each other, and drop from the total integration. Therefore

$$h(z) = \frac{1}{2\pi i} \int_{\Gamma} \frac{F(z)}{z-a} dz \quad (30)$$

as in eq. (26). As before, $F_p(z) = F(z) - h(z)$. Finally, the source term inside D is added back, as in eq. (27). The boundary value of F_p obtained after this source term is added back may be multi-valued, with a branch cut discontinuity in the imaginary part that is a multiple of $\pi/2$. This is a valid input to the next step, as explained in section 3.6 below (following eq. (31)).

The choice of the added singularities' locations outside D (inside the holes) does not affect the results: changing the location of a singularity with some charge from z_a to z_b is equivalent to adding a singularity with opposite charge at z_a , and with the same charge at z_b . Since these two singularities lay outside Ω , and are of opposite charge, this amounts to adding an analytic function to F , which is removed by the Cauchy integral technique described above. The result is also unaffected by the location z_{in} of the source term subtracted inside D , since its location only affects the F_p and it is later added back. The choice of cuts in eq. 29, of course, does not affect the Cauchy integral calculation, since the cuts do not enter the final calculation, eq. (30).

3.6 Calculating the sources' locations

At this point, we are given the value of $F_p(z)$ on the boundary Γ of the domain D . Using eq. (12), we define

$$\xi(z) \equiv \exp[4F_p(z)] = \exp\left[\sum_{m=1}^{n_c} k_m \ln(z - z_m)\right]. \quad (31)$$

Note that while F_p may be have a multi-valued imaginary part with a branch cut discontinuity: $i\theta \rightarrow i\theta + in\pi/2$, for some integer n , $\xi(z)$ is single valued, since $\exp[4(in\pi/2)] = 1$.

Since the k -values, k_m , are integers, we can assume without loss of generality that $k_m = \pm 1$. Other charges may be formed by placing several sources at the same location. We can therefore write eq. (12) as:

$$4F_p(z) = \sum_{m=1}^{n_+} \ln(z - z_m^+) - \sum_{m=1}^{n_-} \ln(z - z_m^-), \quad (32)$$

where n_+, n_- are the number of sources with $k_m = +1, -1$ respectively.

Remark 2 According to the divergence theorem, the difference $n_+ - n_-$ is equal to $\Phi_{\Gamma}/(\pi/2)$ (this can also be shown directly by calculating the integral of F_p along Γ). Since Φ_{Γ} is known (see also remark 1), $n_+ - n_-$ is known. Therefore, to fix n_+, n_- only one number remains to be chosen, e.g. $n_+ + n_-$. The choice of the total number of sources affects the accuracy of the reconstruction of the input data, see sections 3.8 and section 5.

Substituting eq. (32), eq. (31) now reads

$$\xi(z) = \frac{\prod_{m=1}^{n_+} (z - z_m^+)}{\prod_{m=1}^{n_-} (z - z_m^-)}. \quad (33)$$

In eq. (33), the numbers z_m^+, z_m^- are the roots of polynomials, and we proceed by presenting the polynomials in a different form. Later, once the polynomials are found, the z_m^+, z_m^- are recovered by finding the roots of the polynomials. We rewrite eq. (33) as

$$\xi(z) = \frac{\sum_{m=0}^{n_+-1} p_m z^m + z^{n_+}}{\sum_{m=0}^{n_--1} q_m z^m + z^{n_-}}, \quad (34)$$

where p_i, q_i are unknown coefficients. Note that the prefactors of z^{n_+}, z^{n_-} are indeed equal to 1, as can be seen by expanding the polynomials in eq. (33). To find the unknown coefficients numerically, we discretize eq. (34), and give $\xi(z)$ at N points $z_j, j = 1..N$, and write $\xi_j \equiv \xi(z_j)$:

$$\xi_j = \frac{\sum_{m=0}^{n_+-1} p_m z_j^m + z_j^{n_+}}{\sum_{m=0}^{n_--1} q_m z_j^m + z_j^{n_-}}. \quad (35)$$

The RHS of eq. (35) is a *rational function interpolation* of the ξ_j data (see e.g. [25]). In the implementation described below, the unknowns p_m, q_m are evaluated as follows. Rearranging, eq. (35) reads

$$\xi_j - z_j^{n_+-n_-} = \sum_{m=0}^{n_+-1} p_m \left(z_j^{m-n_-} \right) - \sum_{m=0}^{n_--1} q_m \left(\xi_j z_j^{m-n_-} \right). \quad (36)$$

These are N linear equations for the $n_+ + n_-$ unknowns: p_m, q_m . Since N will typically be larger than $n_+ + n_-$, the solution will only be approximate, e.g., a solution in the least-mean-square (LMS) sense.

Once the unknown variables p_m, q_m are found, the source terms' locations z_m^-, z_m^+ are calculated by finding the roots of the two polynomials appearing in the RHS of eq. (35).

Remark 3 *A solution in the LMS sense, as described in eq. (36), is not recommended if the resulting errors in the ξ_j 's are large, when compared with ξ_j . This is due to two reasons: First, the input data ξ_j appears in both sides of eq. (36), so the errors in the LMS solution might not reflect the errors in ξ_j , and moreover, the LMS solution may be sensitive to the choice of the coordinate system origin. Secondly, an error in ξ_j scales as a difference in the resolution (since $|\xi_j| = \exp(\phi)$), rather than the more "natural" error definition given by the ratio of input and obtained resolution. This means that at low resolutions (large cell-sizes), errors in ξ might represent large relative resolution deviations.*

3.7 Restoring the harmonic part

Once the sources' locations have been determined, the function ϕ can be calculated everywhere in the domain. The solution will usually be approximate, i.e., the resulting sum of sources will only approximate the required $F_p = F - h$, and some deviation from both boundary alignment and resolution requirements may be found. The trade-off between the two requirements can be partly controlled by choosing an appropriate harmonic ϕ_L in eq. (6). In order to satisfy the cell-size requirements exactly, one can solve

the (direct) Poisson equation, given by eq. (6), but this time with a known charge distribution, and with Dirichlet boundary conditions, i.e. by specifying ϕ on the boundary.

A different approach is to try and satisfy the boundary alignment conditions, by solving eq. (6) with Neumann boundary conditions. This gives a new ϕ on the boundary. The difference between the input (required) ϕ , and the ϕ obtained on the boundary can be used to evaluate the quality of the source distribution calculated by the IP algorithm, and the total number $n_+ + n_-$ of sources can be changed accordingly. Note that when D is not simply connected, solving ϕ_L with Neumann boundary conditions (Condition 2) does *not* mean that exact boundary alignment is obtained, because Condition 4 might not be exactly fulfilled.

3.8 Summary of the IP algorithm

The steps followed in calculating the source distribution (locations and charges) can be summarized as follows:

1. Adjusting the input resolution near the junction points (section 3.4.1).
2. Subtracting the junction source terms (sections 3.4.2, 3.5).

$$F \rightarrow F - \sum_{\{junctions\}} \frac{Q_{J_i}}{2\pi} \ln(z - z_j^J) + \sum_{m=2}^{n_{bnd}} \sum_{J_i \in \Gamma_m} \frac{Q_{J_i}}{2\pi} \ln(z - z_m^{hole}) - \frac{k_{in}}{4} \ln(z - z_{in}) - \sum_{m=2}^{n_{bnd}} \frac{k_{hole_m}}{4} \ln(z - z_m^{hole}).$$

3. Calculating the Cauchy integral to find $h(z)$ on the boundary. Finding $F_p = F - h$ on the boundary. (section 3.5).
4. Adding back the source inside D (section 3.5):

$$F_p \rightarrow F_p + \frac{k_{in}}{4} \ln(z - z_{in}).$$

5. Calculating the locations of sources (section 3.6), using some initial total number of sources $n_+ + n_-$.
6. Solving the (direct) Poisson problem to find the new F approximating the original F . According to the quality of reconstruction, step 5 can be repeated with a different total number of charges.

4 Creating the final mesh

Once the function ϕ is set throughout the domain, the final mesh can be constructed. There are various possible approaches to this problem. We present a simple method that was used to create the examples in section 6. In this method, the domain is cut along geodesics that follow the cross-field directions. Every source is at the end of a cut, and every inner boundary component is connected to the outer boundary by a cut, see fig. 5, a.. The direction of the geodesics emanating from the sources and directed along the cross-field are calculated as explained in Appendix A. The cut domain, which we denote by D' , does not contain any sources in its interior, and its boundary consists of a single connected component. A conformal mapping g of D' into the plane can be constructed, with a conformal factor (local scaling) of $\exp(\phi)$.

We consider D' and $g(D')$ as lying in the complex plane. First, note that since D' contains no sources in its interior and is simply-connected, $\theta(z)$ is uniquely defined in D' . Denote by $\theta_{D'}(z)$ this single-valued $\theta(z)$. To calculate the mapped boundary $g(D')$, we calculate the function $g(z)$ along the (single) boundary Γ' of D' :

$$g(z)|_{\Gamma'} = \int_{\Gamma'} \exp(\phi(z) - i\theta_{D'}(z)) dz.$$

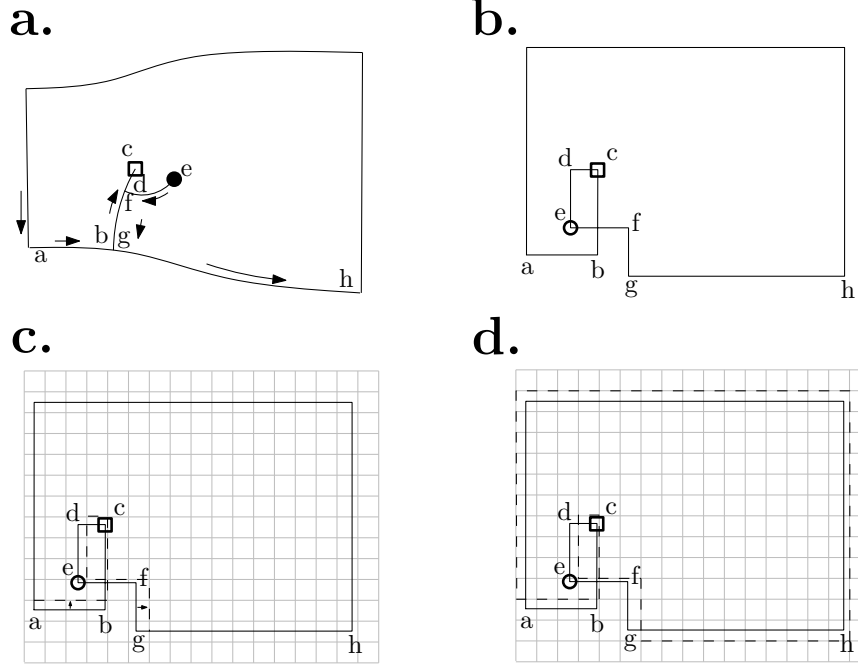


Figure 5: Creating the final mesh. **a.** The cut-tree. Points a, c are source locations. **b.** The mapped cut-tree. **c.** Modifying the boundary. Integer grid lines are drawn in gray. Sections $\overline{ab}, \overline{bc}, \overline{cd}, \overline{de}$ and \overline{ef} are shifted to the nearest grid line (dashed line). Section \overline{fg} is shifted so that \overline{de} and \overline{ef} will be of equal length. **d.** Section \overline{gh} is shifted so that the lengths $\overline{bc} = \overline{cd} + \overline{fg}$. The process is continued until point a is reached again.

If this mapping is used to create a mesh, by placing a Cartesian grid on $g(D')$, and mapping it back to D' by taking the inverse of g , the mesh obtained will be invalid: the grid will not be continuous across the cuts, will contain invalid cells at the sources, and cut cells at the boundaries. To correct this, and obtain a valid final mesh, we seek a function $f(z)$ on Γ' approximating $g(\Gamma')$. A related problem, for surface meshes without boundary alignment, was addressed in [5].

First, note that because the cuts are added along geodesics aligned with the cross field, $g(\Gamma')$ is a piecewise straight line, whose straight segments are either horizontal or vertical. In order to obtain a valid final mesh, the new function $f(z)$ should have the following properties:

1. $f(z)$ must map the closed path Γ' to a closed path.
2. $f(\Gamma')$ must be a piecewise straight line, composed of horizontal and vertical lines only. To each segment in $g(\Gamma')$ corresponds a segment in $f(\Gamma')$, at the same order along the paths.
3. Each segment in $f(\Gamma')$ must be directed in the same direction as the corresponding segment in $g(\Gamma')$, or be of zero length.
4. If γ' is a geodesic cut of Γ' , it is followed in both directions: γ' and γ'^{-} . Then $f(\gamma')$ and $f(\gamma'^{-})$ must have the same length. This ensures the same number of cells on both sides of a single cut.

We describe a simple algorithm for creating $f(\Gamma')$, meeting these requirements. To create $f(\Gamma')$, the path $g(\Gamma')$ is followed, and the segments of $g(z)|_{\Gamma'}$ are modified one after the other, in the order they appear in $g(\Gamma')$, see fig. 5. Each segment is shifted to a parallel line of the integer grid. This is usually the nearest segment, except in the case the length of the segment directly before this segment dictates a different shift, according to properties (3,4) above, see fig. 5,c.. This algorithm requires that at least one junction exist in Γ , which is mapped to a right angle in $g(\Gamma')$. The algorithm starts from the segment following this junction, so that when the last segment is reached and shifted, a change in the length of the first segment will be allowed.

The resulting modified path is similar to $g(\Gamma')$, only with different segments lengths. $f(\Gamma')$ is then defined as the composition of $g(\Gamma')$ followed by a linear mapping of each segment of $g(\Gamma')$ to the corresponding segment in the modified path. Once $f(\Gamma')$ is known, $f(D')$ can be defined. In the implementation described below this is done by placing a triangular mesh in D' , and solving the Laplace equation for both $\text{Re}(f(z))$ and $\text{Im}(f(z))$ on D' . $f(z)$ obtained in this way is not in general analytic, since the CR conditions might not hold between $\text{Re}(f(z))$ and $\text{Im}(f(z))$, so the final mesh will not exactly conformal, but only approximately, as $f(z)$ approximates the analytical function $g(z)$. The mesh edges are then extracted by tracing the integer-valued lines of $\text{Re}(f(z))$ and $\text{Im}(f(z))$.

Once the connectivity of the mesh has been established, a final smoothing procedure which involves all the interior vertices, including those on the cut path, is preformed. In the definition of $f(z)$, both $\text{Re}(f(z))$ and $\text{Im}(f(z))$ are harmonic functions. A standard finite difference approximation to a harmonic function on a uniform grid, assigns to each vertex the average of its four neighboring vertices. This is applied to all vertices inside the mesh, including those on the cuts of Γ' , which can be interpreted as a generalization of the definition of $f(z)$ (to be more precise, in this smoothing procedure it is $f^{-1}(z)$ which is assumed to be harmonic). This procedure is the well-known Laplacian smoothing procedure, in which the location of every vertex is equal to the average location of its neighboring vertices, and in this application it is theoretically justified if $f(z)$ is approximately $g(z)$, i.e. if $f(\Gamma')$ is close to $g(\Gamma')$.

5 Implementation Details

The algorithm described in the previous section was implemented in Matlab [27], using standard Matlab functions, such as a least-mean square solver for linear systems, an ODE solver, etc..

To calculate the Cauchy integral (section 3.5), the boundary was approximated by a piecewise linear path between sampled points, with a constant value of ϕ on each segment. This allows the integral over each segment of the path to be calculated analytically, and the path integral there is equal to the sum

of the integral over the segments. This is a low order approximation, but was sufficient for the present purposes.

The linear system, eq. (36), was solved in the least-mean-square sense, using Matlab's `mldivide` function. Note that the coefficients in this system contain powers of the z_j 's, the boundary data locations. This leads to ill-conditioned linear systems, which are more sensitive to the numerical round-off errors, as the number of sources n_+, n_- increases, or the range of $|z|$ increases. Tests indicate that when the solver issues a warning, indicating that the matrix is rank deficient to the working precision, the results of the IP algorithm become unreliable (this was not the case in the examples in section 6). For input domains where many sources or a high $|z|$ range are required, a different solver implementation may be needed, such as a varying precision arithmetic computation, working at higher floating-point accuracies.

After sources' locations are recovered by taking the roots of the polynomials, the function ϕ_L of eq. (6) is calculated, given the Neumann boundary conditions in eq. (8). This was done using the Method of Fundamental Solutions (MFS), see e.g. [26], where the solution to the Laplace for ϕ_L is approximated by a sum of (real) source terms, i.e. functions of the form $\phi(r) = \frac{Q_i}{2\pi} \ln |\vec{r} - \vec{r}_i|$, with the locations \vec{r}_i lying *outside* the domain D where ϕ_L is harmonic. The charges Q_i are then calculated to best fit the boundary conditions.

Once the function ϕ is given everywhere in D , the directions of the star-geodesics at each source are calculated, as described in Appendix A. The star-geodesics, emanating from the source were then traced by solving the differential equation, eq. (2), using a Runge-Kutta ODE solver. The cut-tree was constructed by taking one source at a time at some arbitrary order, and introducing a cut along a star-geodesic of the added source, that is closest to the cut-tree constructed so-far. This construction of the cut tree is admittedly arbitrary, and may not be optimal in some cases.

Lastly, the algorithm described in section 4 is applied.

6 Example Results

In the first example, a domain with two boundary components is meshed, see fig. 6.a. The side lengths of the outer and inner squares are 8 and 2.26 respectively, centered at the origin, and the inner square is rotated at angle $\alpha = 0.35\pi$ radians. The ratio of input resolution between outer and inner boundaries is 2.5. The inner and outer boundaries were sampled with 268 and 76 points, respectively, which were used both for calculating the Cauchy integral, and for the sources' locations calculation. Singularities of ϕ are not required at the junctions points, since all inner angles are multiples of $\pi/2$ radians. As there is no flux of $\nabla\phi$ through either boundary (because the boundaries' curvature is $\kappa = 0$, see eq. (6), and ϕ is regular at the junctions), the total charge inside is zero (see remark 2). Hence an equal number of $k = +1$ and $k = -1$ charges were used. Fig. 6.a., shows the domain boundary, together with the calculated sources' locations. Squares mark sources with $k = +1$, and circles sources with $k = -1$. The distribution is composed of 18 sources of each type, i.e. a total of 36 sources. The value of ϕ obtained by solving the Poisson equation with this charge distribution is shown in fig. 6.b., and compared to the input ϕ requested, denoted by ϕ_{input} . The difference $\Delta\phi \equiv \phi - \phi_{input}$ is also plotted. For visual clarity, the average of the ϕ_{input} was subtracted from both ϕ and ϕ_{input} (this is just a scaling of the resolution by a constant factor); this was also done in the resolution comparisons below, in fig. (7).c., fig. (8).b. and fig. (9).c.. The number of sources was chosen to be the smallest for which $\Delta\phi < 0.1$ at all boundary points. This criterion is used for choosing the number of sources is also used in the following examples. The cut-tree used in creating the final mesh is shown in Fig. 6.c., and the final mesh in shown Fig. 6.d.

The domain in the second example is the union of three unit radius circles which pass through the origin, see fig. 7.a. The boundary has three junctions, all with the same inner angle, $\theta_{in} = 5\pi/3$. Boundary alignment requires ϕ to be singular at the junctions, corresponding to sources with charge $Q = -\pi/5$, or k -value $k = Q/(\pi/2) = -2/5$, see eq. (17),(18). As explained in section 3.4.1, singular functions can be added the input resolution in order to avoid clustering of sources near the junction points. In this example, we compose the input ϕ from a sum of pairs of source terms, with k -values as marked in fig. 7.a.. The pairs of nearby sources with opposite charge create the desired singularities at the junctions, and have only a small effect at distances large compared with the distance between the sources of each pair.

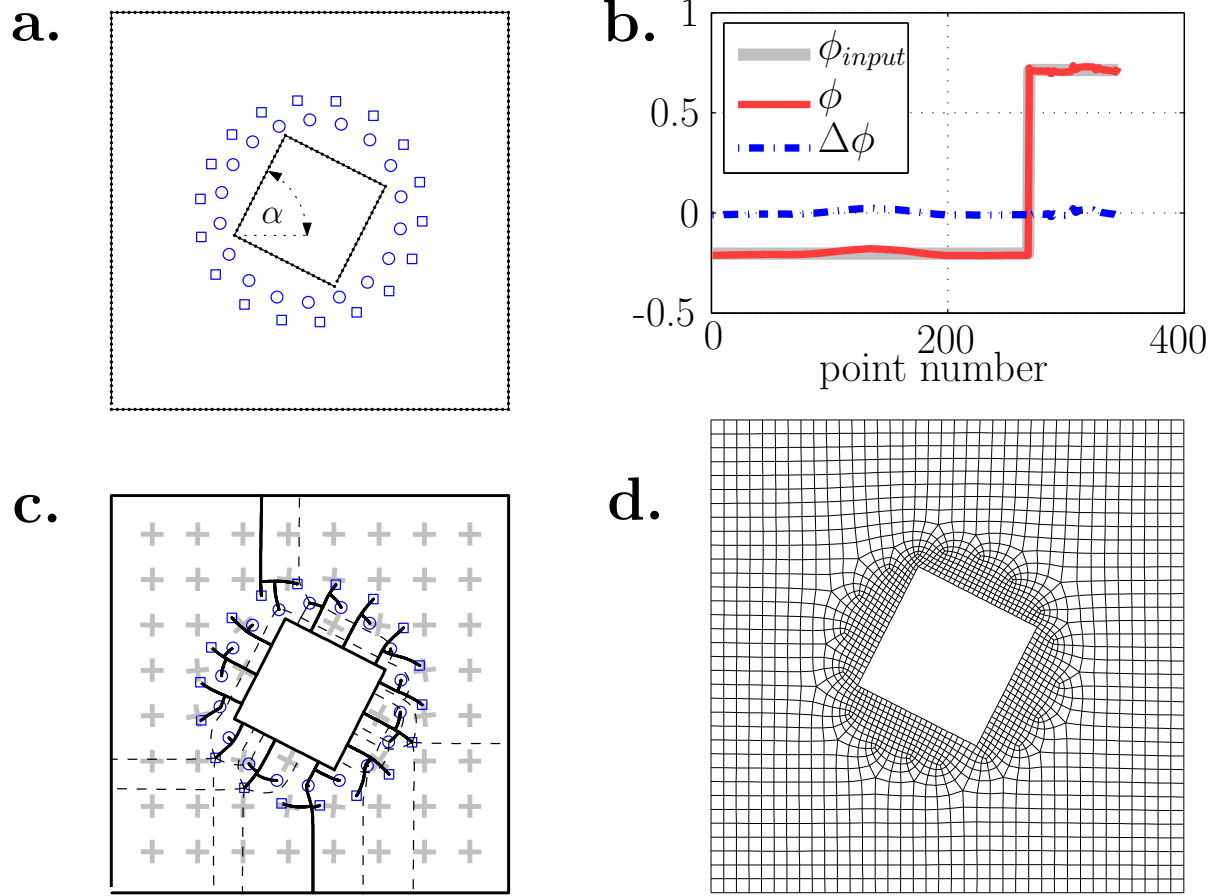


Figure 6: (Color online) Mesh generation procedure: **a.** Domain boundary (solid line) sampled at points (dots) as input to the IP algorithm. Resulting charge distribution is plotted: $k = +1, -1$ cone-points are plotted with squares and circles, respectively. **b.** ϕ obtained at boundaries. Points 1-268 correspond to the outer boundary component, the rest to the inner component. Input ϕ , obtained ϕ and their difference are plotted. **c.** Crosses at selected points (crosses), cut-tree composed of the boundary and star-geodesics (solid line), and selected additional star-geodesics (dashed lines). **d.** The final mesh.

We can now calculate the total flux of ϕ through the boundary in this example. The (singular) flux through each junction is $\pi/6$ (positive value when $\nabla\phi$ is directed outward). The flux through each of the three arcs is $\int \partial_n \phi ds = -\int \partial_s \theta ds = -\Delta\theta$ along the arc, and since for each arc $\Delta\theta = 4\pi/3$, the total flux through the boundary is $3 \cdot (\pi/6) + 3 \cdot (-4\pi/3) = -7(\pi/2)$. This is in agreement with the formula given in remark 1. Therefore, the sum of k -values of the sources inside the domain should be -7 , so $n_- = n_+ + 7$. Hence, k_{in} of eq. (23) was -7 ; $z_{in} = 1$ was used. The lowest number of sources for which $\Delta\phi < 0.1$, shown in fig. 7,b., has $n_+ = 2$, $n_- = 9$. The boundary was sampled at 525 points.

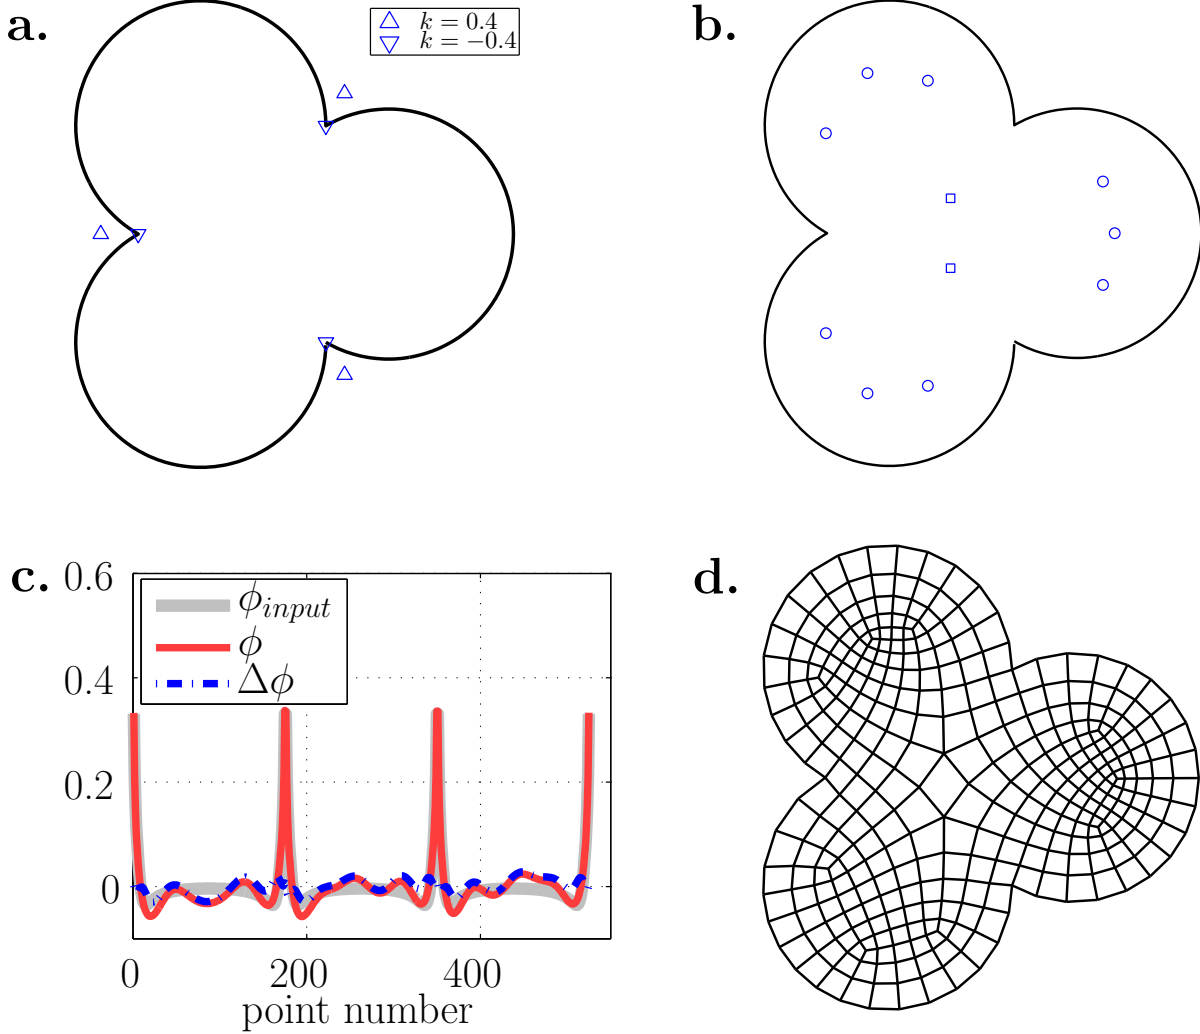


Figure 7: (Color online) **a.** Boundary with adjustment of resolution near junction points. **b.** Obtained source locations (blue). **c.** Input ϕ , obtained ϕ and their difference. **d.** Final mesh.

The boundary in the third example is composed of two boundary components, each with a single junction point. The inner angles at the inner and outer boundary junctions are $\theta_{in} = 4\pi/3, \pi/2$ respectively, so a singularity of ϕ will form at the inner boundary junction. As in the previous example, we add a singular function to the input ϕ , by adding a pair of source terms, one at the junction point with $k = 1/2$, as obtained from eq. (17),(18) with $\theta_{in} = 4\pi/3$. No inner source term is required, and from the flux trough the hole boundary it follows that $k_{hole} = -3$. After subtracting the junction source term with $k = 1/2$ (see section 3.4) a source with total charge of $k_{hole} = -3 - 1/2$ is added inside the hole, see fig. 8,a.. The inner and outer boundary components are sampled at 325 and 142 points, respectively. The distribution with $n_+ = n_- = 16$, the smallest number of sources for which $\Delta\phi < 0.1$, is plotted at

fig. 8,a., with squares and circles representing charges as in previous examples. Two pairs of $k = \pm 1$ charges are located inside the inner boundary, i.e., outside the domain. Their constitute a correction to the homogeneous ϕ_L , which is very small, since in each pair the opposite charges almost overlap. The obtained ϕ reconstruction is given in fig. 8,b.. The final mesh, fig. 8,c., contains 26 irregular vertices, 13 of each type. This is because two pairs of sources lay outside the domain, and another pair (the pair of sources closest to the inner junction) was “eliminated” in the creation of the final mesh (section 4): the two opposite charges where united, giving a zero total charge.

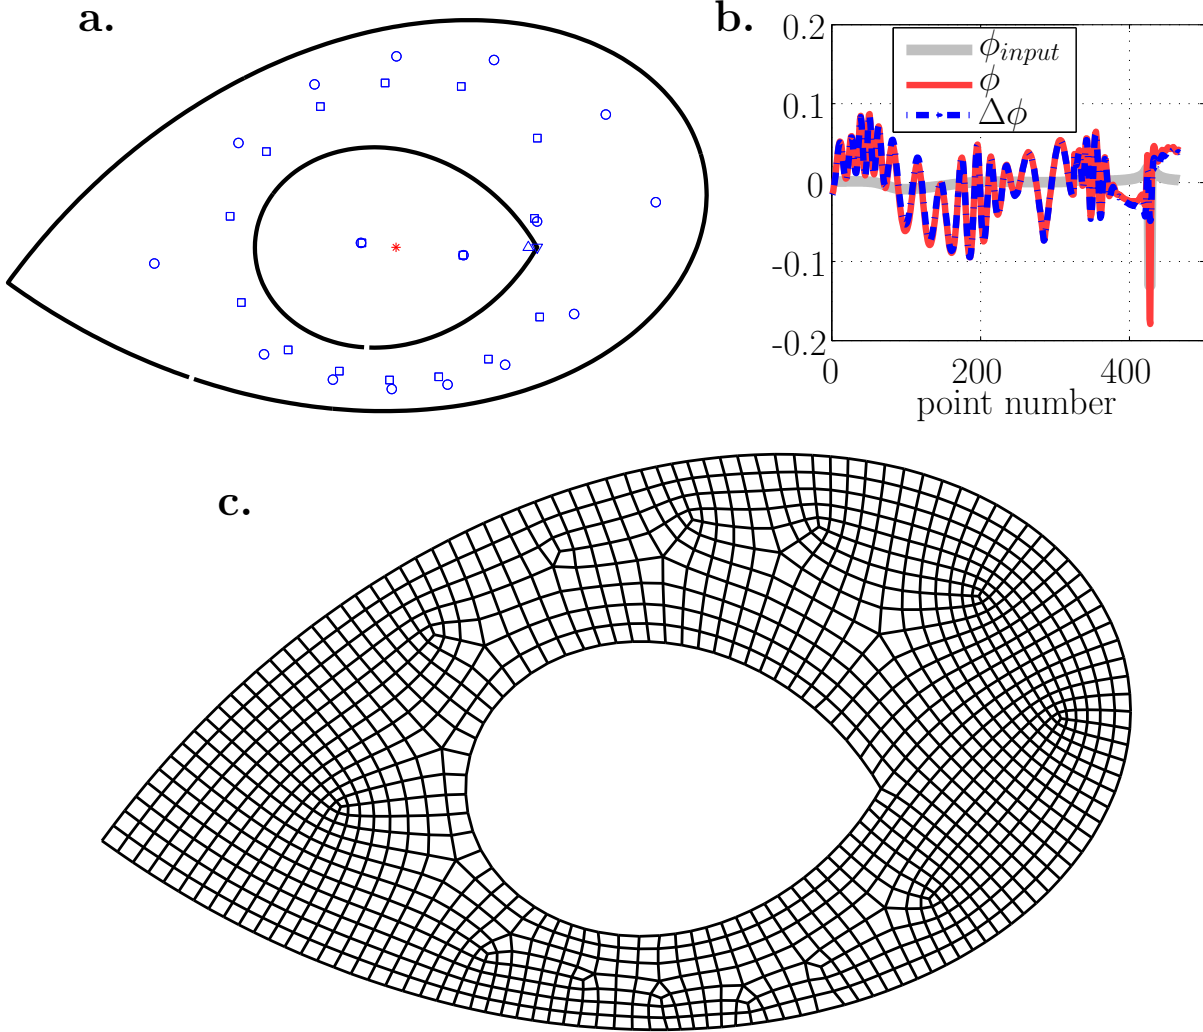


Figure 8: (Color online) **a.** The boundary (solid line), sources to adjust input resolution near junction (triangles), source distribution (squares and circles), and inside source location (red star), used in removing the harmonic part. **b.** Input ϕ , obtained ϕ and their difference. **c.** Final mesh.

In the final example the boundary of the square $[-2, 2] \times [-2, 2]$ was assigned a varying input resolution proportional to:

$$resolution \propto 1.5 + (x + 2)^2 + \frac{(y + 2)^2}{2}.$$

The resolution thus varies by a factor of 17 within the domain. As always, eq. (11), $\phi = \ln(resolution)$. The total charge is zero, so $n_+ = n_-$. The boundary was sampled at 536 points. A source distribution with $\Delta\phi < 0.1$ was obtained with just 8 charges, i.e. with $n_+ = n_- = 4$. Fig. 9,a. shows the source distribution, Fig. 9,b. shows the obtained ϕ vs. the input ϕ , and the final mesh is shown in Fig. 9,c..

It is instructive to compare the star-geodesics and cut-tree with the final mesh. The two are overlaid in Fig. 9,d. (this would create a “crowded” appearance when more sources are present). The black solid line traces the cut-tree, the dashed lines are additional star-geodesics, and the mesh is drawn with the background in gray lines.

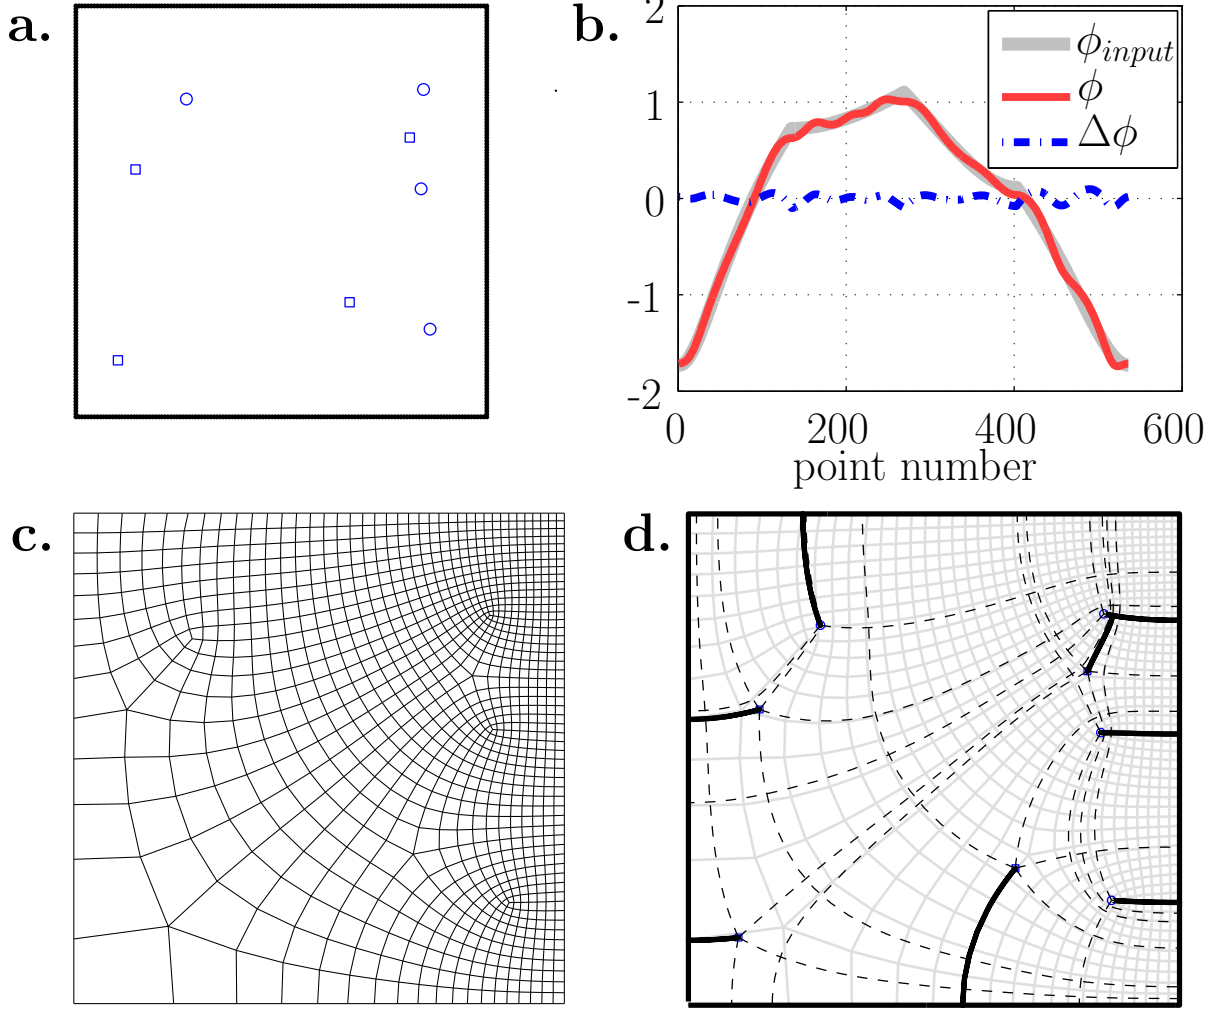


Figure 9: (Color online) **a.** Boundary, source distribution (squares and circles). **b.** Input ϕ , obtained ϕ and their difference. **c.** Final mesh. **d.** Final mesh (gray lines), overlaid with the cut-tree (solid line) and additional star-geodesics (dashed lines).

Cell shape quality was measured using a variant of the quality measure β [28], as defined in [29]. $\beta = 1$ represents a square cell, while $\beta = 0$ represents a cell with an inner angle of π . Cells with a high aspect ratio are also given low β -values. Table 1 shows the minimum and average β values, and the total number of cells, for the examples in Fig. 6-9. For each example, the tabulated information is given for the mesh shown in the corresponding figure, and for a finer mesh of the same domain, where the resolution function was doubled (producing a mesh with about four times the number of cells). Note that multiplying the resolution by a factor amounts to adding a constant to ϕ , which does not change the source distribution obtained from the IP algorithm, as this constant, which is part of $h(z)$, is readily removed when $h(z)$ is subtracted.

	num. cells	min. β	avg. β
Fig. 6	1898	0.576	0.945
	7630	0.627	0.975
Fig. 7	316	0.603	0.901
	1036	0.620	0.942
Fig. 8	1233	0.506	0.933
	4895	0.562	0.958
Fig. 9	1062	0.626	0.964
	6059	0.684	0.987

Table 1: Mesh statistics.

7 Conclusions and Future work

An unstructured quadrilateral mesh generation scheme in planar domains was presented. The method rests on a theoretical foundation, linking the mesh generation problem with the Inverse Poisson (IP) problem. An IP solution algorithm is presented, whose output is interpreted as the location and type (degree) of irregular vertices in the domain. The continuum fields obtained, describing mesh resolution and directionality, are conformal everywhere except on the irregular vertices, and fit the required input properties at the boundary, or at other user-defined locations. An algorithm for creating a valid final mesh is also presented. Example meshes feature irregular vertices where they are needed, in combination with highly regular regions where possible.

Directions for future work include more sophisticated methods for solving the rational function interpolation equations, and for constructing the final mesh. The relations between conformal unstructured mesh generation, the IP problem and rational function interpolation raise many research questions. These may lead to a deeper understanding of the properties of high quality meshes, and to better algorithms for creating them.

Acknowledgements. The author would like to thank Mirela Ben-Chen, Shlomi Hillel, Dov Levine, Yair Shokef and Vincenzo Vitelli for helpful discussions and critical reading of the manuscript.

8 Appendix A: Star-geodesic directions

A cross is defined at every point p of the domain, that is not a singularity of the function ϕ . Geodesic curves that start at p can be drawn in all four directions of the cross. By the definition of a cross-field, such a geodesic will be aligned with the cross-field everywhere along the curve. If p is a singular point, with k -value $k_p \neq 0$, a cross is not defined at p , but there are $4 + k_p$ geodesics that are incident on p and follow the cross-field directions elsewhere on the curve. These will be called *star-geodesics*. For example, the geodesics drawn in fig. 6,c. and 9,d. are star-geodesics.

Denote the angle from the x -axis around p by ψ , and the cross direction when p is approached from direction ψ by $\theta(\psi)$, see fig. 10. To calculate the directions in which star-geodesics emanate from p , we first calculate the cross when p is approached from some direction, e.g. the positive x -axis. This can be done using eq. (3) along a curve from the boundary to p which approaches p from the direction $\psi = 0$. Once $\theta(0)$ is known, $\theta(\psi)$ for any ψ can be calculated by using eq. (3) along a small circular arc around p at radius r , α_r . The singularity term at ϕ is $\frac{k}{4} \ln |r|$, and according to eq. (3)

$$\theta(\psi) = \theta(0) + \int_{\alpha_r} \frac{\partial \phi}{\partial n} ds = \theta(0) + \frac{\partial \phi}{\partial r} \psi r = \theta(0) + \frac{k}{4r} \psi r = \theta(0) + \frac{k}{4} \psi. \quad (37)$$

Note that contributions to ϕ that are regular at p do not affect $\theta(\psi)$ for $r \rightarrow 0$.

The star-geodesic directions are those for which the cross is directed along the ray from p , or

$$\theta(\psi) = \psi + n \frac{\pi}{2}, \quad (38)$$




Figure 10: Calculating the star-geodesics' directions.

with $n \in \mathbb{Z}$. Substituting eq. (37) into eq (38) and rearranging, we find

$$\psi = \frac{\theta(0) - n\frac{\pi}{2}}{1 - \frac{k}{4}}.$$

It is easy to show that there are exactly $4 + k$ different ψ -values for which this equation is fulfilled. They are equally distributed around p .

References

- [1] S. J. Owen, A survey of unstructured mesh generation technology, *Proceedings of the 7th International Meshing Roundtable* (1998).
- [2] G. Bunin, A continuum theory for unstructured mesh generation in two dimensions, *Computer Aided Geometric Design* **25**, pp. 14–40 (2008).
- [3] P. Alliez, M. Attene, C. Gotsman and G. Ucelli, Recent Advances in Remeshing of Surfaces, in *Shape Analysis and Structuring*, L. De Floriani and M. Spagnuolo (Eds.) (Springer, 2007).
- [4] N. Ray, W. C. Li, B. Levy, A. Sheffer and P. Alliez, Periodic Global Parameterization, *ACM Transactions on Graphics* **25**(4) (2006).
- [5] F. Kälberer, M. Nieser and K. Polthier, QuadCover - Surface Parameterization using Branched Coverings, *Computer Graphics Forum*, **26**(3), pp. 375–384 (2007).
- [6] K. Shimada, J. Liao and T. Itoh, Quadrilateral Meshing with Directionality Control through the Packing of Square Cells, *Proceedings of the 7th International Meshing Roundtable*, pp. 61–76 (1998).
- [7] N. Ray, B. Vallet, W. Li and B. Levy, N-symmetry direction fields on surfaces of arbitrary genus, *Tech Report* (2006).
- [8] X. Gu and S. Yau, Global Conformal Surface Parameterization, *Eurographics Symposium on Geometry Processing* (2003).
- [9] T. Aubin, *Some nonlinear problems in Riemannian Geometry* (Springer-Verlag, 1998).
- [10] S. A. Chang, *Non-linear elliptic equations in conformal geometry* (European Mathematical Society, 2004).
- [11] M. O. Katanaev, *Geometric Theory of Defects*, Phys. Usp. 48, pp. 675–701 (2005).

- [12] L. Kharevych, B. Springborn and P. Schröder, Discrete conformal mappings via circle patterns, *ACM Transactionson Graphics*, **25**(2) (2006).
- [13] M. Yamaguti et al., eds. *Inverse Problems in Engineering Sciences* (Springer-Verlag, Tokyo, 1991).
- [14] M. Hämmäläinen, R. Hari, R. J. Ilmoniemi, J. Knuutila and O. V. Lounasmaa, Magnetoencephalography - theory, instrumentation, and applications to noninvasive studies of the working human brain, *Reviews of Modern Physics* **65** Issue 2 (1993) 413-497.
- [15] A. A. Ioannides et. al, Continuous probabilistic solutions to the biomagnetic inverse problem, *Inverse Problems* **6**, pp. 523-542 (1990).
- [16] P. Johnston, ed., *Computational Inverse Problems in Electrocardiography* (Southampton: WIT Press, 2001).
- [17] D. Zidarov, *Inverse Gravimetric Problem in Geoprospecting and Geodesy* (Amsterdam: Elsevier, 1990).
- [18] T. Ohe and K. Ohnaka, Precise Estimation Method for Locations in an Inverse Logarithmic Potential Problem for Point Mass Models, *Applied Mathematical Modelling*, **18**(8), pp. 446-452 (1994).
- [19] A. El-Badia and T. Ha-Duong, An inverse source problem in potential analysis, *Inverse Problems* **16**, pp. 651-63 (2000).
- [20] T. Nara and S. Ando, A projective method for an inverse source problem of the Poisson equation, *Inverse Problems*, **19**(2), pp. 355-369 (2003).
- [21] L. Baratchart, A. Ben Abda, F. Ben Hassen and J. Leblond, Recovery of pointwise sources or small inclusions in 2D domains and rational approximation, *Inverse Problems* **21**, p. 51-74 (2005).
- [22] J. B. Conway, *Functions of One Complex Variable I* (Springer, 1997).
- [23] B. Noble, *Methods based on the Wiener-Hopf technique for the solution of partial differential equations* (Pergamon Press, 1958).
- [24] S. Axler, P. Bourdon and W. Ramey, *Harmonic Function Theory*, 2nd edition (Springer, 2001).
- [25] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 2nd edition (Cambridge University Press, 1992).
- [26] G. Fairweather and A. Karageorghis, The method of fundamental solutions for elliptic boundary value problems, *Adv. Comput. Math.* **9**, pp. 69-95 (1998).
- [27] MATLAB, The Mathworks Inc..
- [28] S. H. Lo, Generating quadrilateral elements on plane and over curved surfaces, *Comput. Struct.*, **31**, pp. 421-426 (1989).
- [29] S. J. Owen, M. L. Staten, S. A. Canann and S. Saigal, Advancing Front Quad Meshing Using Local Triangle Transformations, *Proceedings of the 7th International Meshing Roundtable* (1998).